Demonstrating the Benefits of Virtual Interface Architecture

A Cooperative Effort Makes 16-node SHV Server Clusters a Reality

VI Architecture: A Brief History

Standard high-volume (SHV) servers have made steady headway in enterprise computing for the last decade. Starting out as departmental servers, Intel Architecture-based SHV servers then progressed to fill the role of mid-range servers and later high-performance servers. Each new generation of Intel processors adds more power, scalability, availability, and price/performance advantages. With the recent introduction of the Intel® Pentium® II Xeon™ processor, SHV servers are serious contenders for high-end business-critical computing in the data center.

However, faster processors from Intel can only take SHV servers so far. Pooling multiple processors within a single server boosts performance higher. But linking multiple servers together, a concept called clustering, is where real performance breakthroughs lie. The idea is to be able to scale a clustered computer system indefinitely by continuously adding more servers and increasing performance proportionate to the number of CPUs added. To do that requires more than very fast server nodes; it also requires very fast interconnections between nodes. This is where Virtual Interface (VI) Architecture comes in.

In 1996, Intel and other companies spearheaded an effort to create a common SHV server interconnect technology that would permit volumes of data to pass very quickly between clustered SHV servers. A core group of companies wrote the VI Architecture specification in 1996, submitted it for industry review in 1997, got buy-in from more than 130 companies in early 1998, and demonstrated VI Architecture-enabled 4- and 6-node clusters throughout 1998.

Now, in September 1998, Intel and other leading technology companies are proud to unveil the latest VI Architecture milestone: a 16-node cluster of SHV servers using an astounding 3 terabytes of storage to store a database of 2 billion records. Clustering SHV servers using VI Architecture is no longer theoretical experimentation. It is a real, commercially available technology that companies can use to solve pressing business problems using common building blocks from any number of leading technology companies.

Single Server Configuration

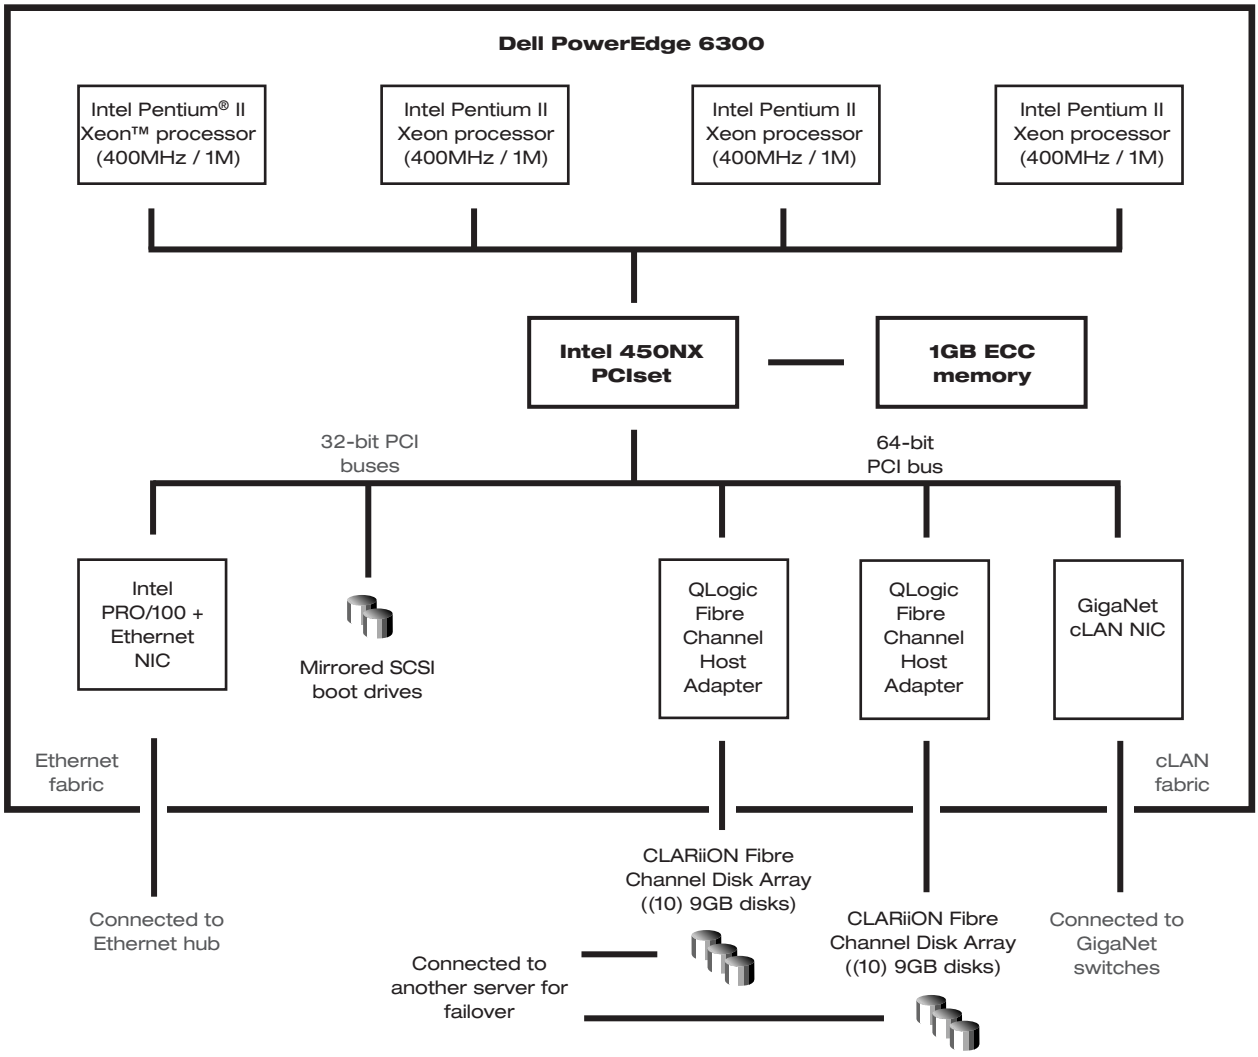

A 16-node Cluster Using VI Architecture

Today's demonstration shows 16 clustered Dell PowerEdge® 6300 quad-processor servers, each containing four Intel Pentium II Xeon processors. Each processor is running at 400 MHz and contains 1MB of L2 cache. Each server is configured with a GigaNet cLAN® GNN1000 network interface card (NIC). These GigaNet NICs are connected together using six 8-port GigaNet cLAN GNX5000 switches. This configuration provides full bandwidth across the entire switch fabric. The GigaNet NICs and switches form a cluster network that allows the servers within the cluster to communicate with each other. The database, IBM's DB2® Universal Database Enterprise Extended Edition v5.2 (DB2 UDB EEE), uses the GigaNet cLAN network to send data back and forth between the servers within the cluster. cLAN is optimized using VI Architecture to provide the lower latency and greater bandwidth required by high-performance clustered database applications.

Each server also contains an Intel PRO/100+ Ethernet adapter. These adapters are all connected to a 24-port Intel Express 220T Stackable Hub. The Ethernet LAN network — which is separate from the GigaNet cluster network described above — is used by the servers to communicate with the clients that access the cluster.

The 2 billion record database is stored on external Fibre Channel disk arrays provided by CLARiiON, the Advanced Storage Division of Data General Corporation. Since DB2 uses a shared-nothing storage model, each server is connected to its own set of disks. To facilitate this connection, each server contains two QLogic* Fibre Channel adapters, each attached to a CLARiiON Fibre Channel-based Disk Array Enclosure (DAE). Each DAE contains 10 CLARiiON 9GB drives, making a total of 20 drives per server and 320 drives for the cluster. The Fibre Channel drives are configured in a mirrored configuration, ensuring that a disk failure cannot bring down the database.

The operating system is Microsoft Windows NT® Server Enterprise Edition version 4.0. Each of the 16 nodes is running its own copy of DB2 and Windows NT Server. The database is storing a total of 2 billion records of demographic data.

Connected to the servers is a Dell Precision Workstation 610 client, also powered by two Intel Pentium II Xeon processors. This workstation is running data visualization software from Visual Insights,® which simplifies the task of looking through millions of records and visualizing results.

Using a web browser, the workstation initiates a query on the database, which causes the data mining algorithms to begin working on the cluster. These algorithms are designed to parse through the database and find

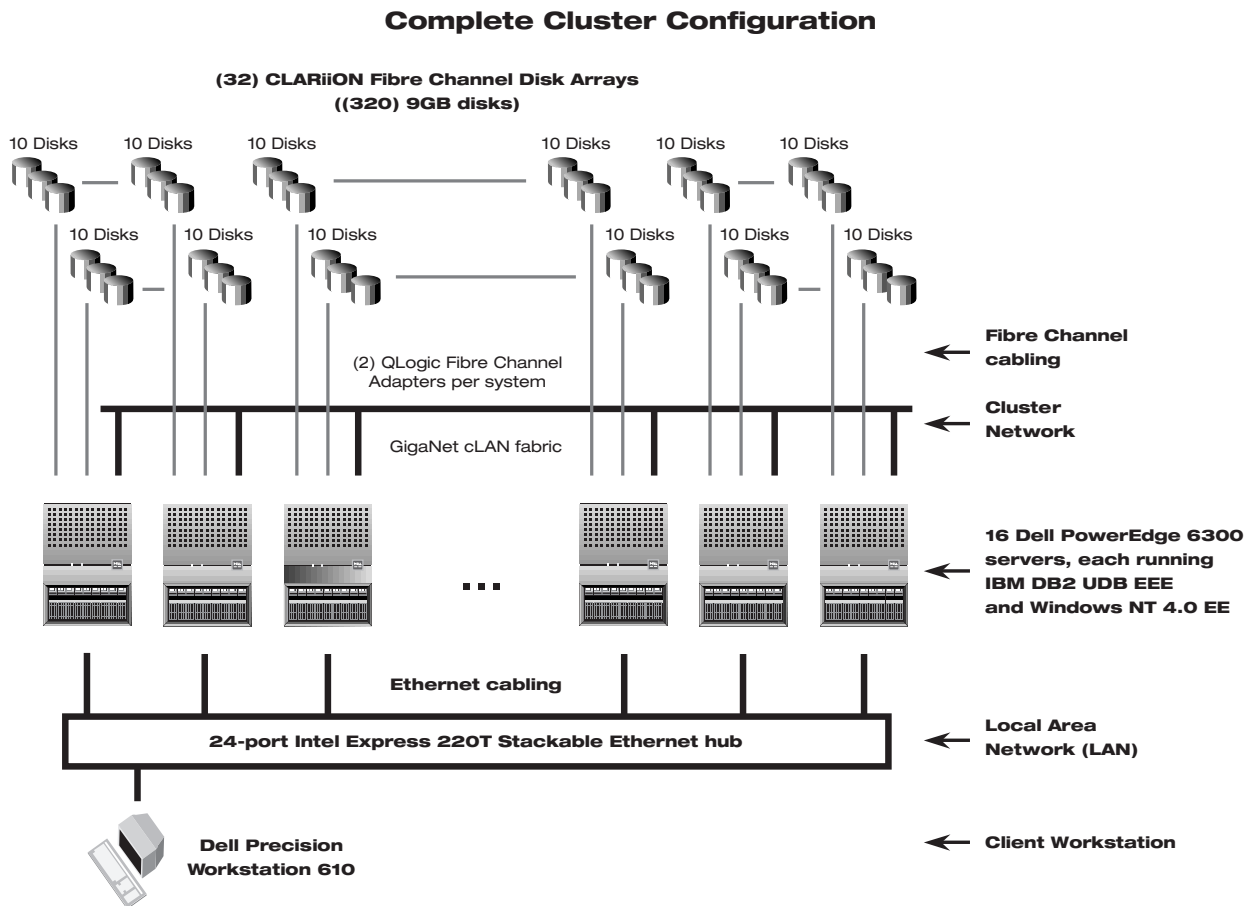

associations hidden within the data. (An example of an association could be that a household with more than three television sets is more likely to own a DVD player than a household with only one or two television sets.) DB2 parallelizes the query both within each 4-way SMP machine, and across all 16 servers in the cluster, making full use of all 64 processors in the cluster. These association tables are fed into the Visual Insights data visualization software and displayed via the web browser. In this demonstration, the Visual Insights software provides an interactive and graphical environment that displays the associations found by the data query, allowing users to explore and validate the relationships in the data set. Visual Insights can also be used to explore the results of the query for patterns and trends that can be used in the decision support activities.

Interoperability of Components

Besides offering the same, or superior, availability and scalability as proprietary clustering solutions, clustered Intel-based SHV servers give you the freedom of interoperability — on two levels.

First, building servers from common building blocks, like off-the-shelf microprocessors and shrink-wrapped operating systems, brings volume economics, simplicity, and continuous innovation. You can choose "best of breed" solutions from multiple vendors for every component in the server. This kind of mix-and-match solution stack is unheard of in the proprietary, monolithic world of legacy computer systems. But it's one of the major advantages of SHV servers.

At the next level, linking individual SHV servers allows you to use low-cost, common server building blocks to build a cluster. It's the logical next step in leveraging the volume economic model pioneered by Intel and the tremendous flexibility and freedom provided by the SHV server model.

CLARiiON

The CLARiiON FC5000 Series Fibre Channel disk array family, featured in the demo, showcases the industry's only full Fibre Channel disk arrays. Fibre Channel is an enabling technology for clustering applications, providing dramatically improved performance, lowering cost of ownership, and delivering the ultimate in high availability. Featuring modular building blocks, parallel deployment, and full end-to-end Fibre Channel technology, the FC5000 Series is a quantum leap in storage technology providing an impressive range of storage solutions for the enterprise, from simple JBOD (just a bunch of disks) configurations to multi-terabyte, highly available RAID storage pools.

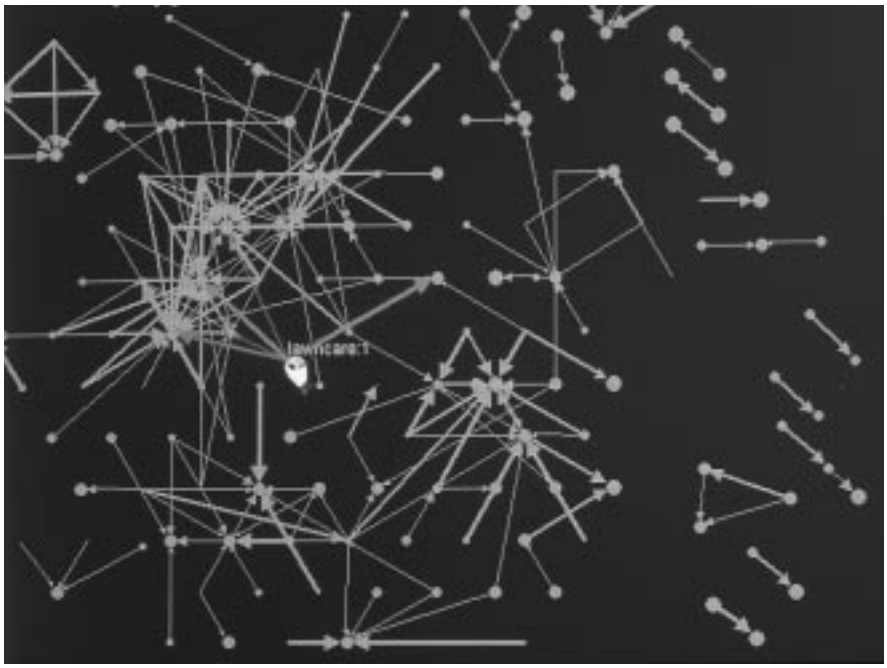

Visual Insights provides the data visualization software that offers a powerful, interactive environment for exploring trends and patterns in the query results provided by DB2.

GigaNet Switch Topology

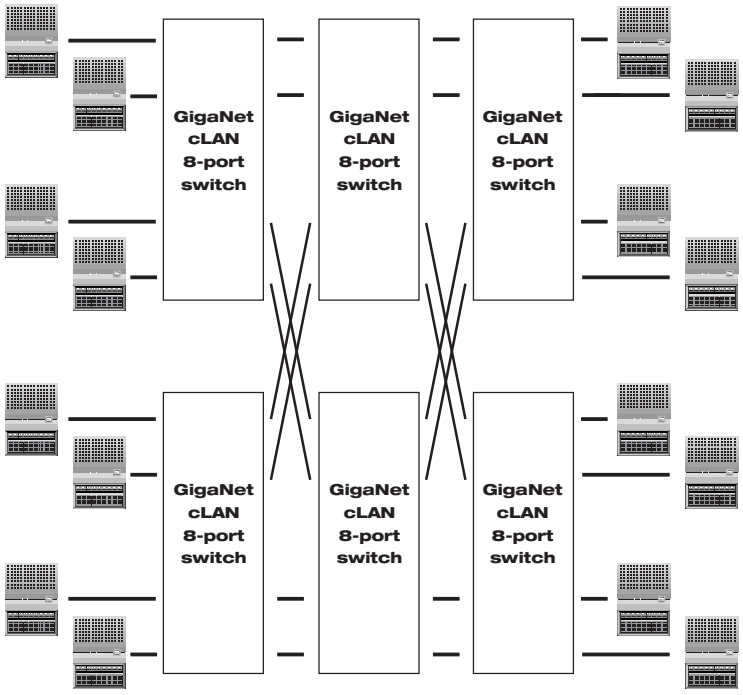

Dell Computer Corporation

The Dell PowerEdge 6300 enterprise-class server features up to four Intel Pentium II Xeon processors, up to 4 GB of ECC RAM, seven PCI slots, and up to 126 GB of internal storage. With the PowerEdge 6300, Dell is bringing its direct sales model to the highest reaches of the enterprise. The simple, cost-effective, plug-and-play cluster model made possible by the PowerEdge 6300 will allow companies with remote locations, like retail stores or banks, to deploy very powerful clusters in these remote locations to provide 24-hour availability without significant maintenance or management demands.

GigaNet Corporation

GigaNet has developed the first native VI Architecture-compliant server interconnects on the market. GigaNet's cLAN products, based on VI Architecture, allow the software applications to exchange messages with almost no networking delay. Throughput rates are much higher, latency is very low, and the CPU overhead is minimal, freeing up the CPU for the primary task of running the application. GigaNet's cLAN family of adapters and switches are based on a low-cost, highly integrated single-chip design. They also have extensive RAS (reliability, availability, scalability) features, needed for mission critical applications. Other features include graceful saturation (does not drop frames or slow performance capacity as traffic levels increase) and auto-configuration. GigaNet cLAN products are uniquely optimized for high availability and scalability applications running on Windows NT Server clusters.

IBM Corporation

IBM's DB2 Universal Database Enterprise-Extended Edition has been used for years in clusters of mainframes and minicomputers. Intel Architecture-based SHV servers are a new and fast-growing realm for this respected big business database. Parallel database technology is often a key component of business intelligence applications typically run on clusters because of the enormous amounts of data to be churned through. It allows you to break the processing into separate execution components that can be run concurrently on multiple processors. The result is dramatically faster response times to queries.

DB2 UDB EEE employs a shared-nothing software architecture in which DB2 server instances run on several machines without the need to directly share memory or disk access. This architecture allows DB2 to support very large databases by dividing the database into partitions that can be stored and managed on separate servers of a shared-nothing hardware system. Because of this unique partitioning architecture and VI Architecture support, DB2 UDB EEE operates well on SHV servers and high-speed network fabrics.

Intel Corporation

The new Intel Pentium II Xeon processor is designed precisely for enterprise server needs like clustering. The 400 MHz core processor speeds, larger L2 cache (up to 2 MB), full-speed L2 cache bus speed, and wider system memory bus all mean more data processed, more web hits handled, more users served, and more business won. The Pentium II Xeon processor, currently running in 4-way configurations, can be clustered to n-way levels of performance — there is no meaningful limit on the number of servers that can be interconnected. Clusters of Pentium II Xeon processor-based servers are ideally suited to enterprise-class, compute-intensive workloads such as data warehousing and data mining. They provide both total performance and price/performance leadership, whether compared with RISC-based solutions or legacy alternatives.

QLogic Corporation

QLogic is a leader in delivering high-performance Fibre Channel interface solutions. Their QLA2100, used in the demonstration, is based on QLogic's award winning ISP2100 PCI-to-Fibre Channel processor. The ISP2100 combines a 64-bit PCI host interface, RISC processor, Fibre Channel protocol engine, dedicated transmit and receive frame buffers, and integrated 100 MB/second Fibre Channel serial transceivers. Three independent DMA channels allow commands and data to be processed simultaneously, and the dedicated frame buffers optimize PCI bus utilization and enable continuous streaming of data during transfer operations. The unique architecture of the QLogic Fibre Channel chip combines low processor overhead and very high throughput to deliver a very efficient interface, resulting in outstanding overall system throughput.

Visual Insights

Visualization is an emerging technology for understanding large, complex, information-rich data sets. Visual Insights, a wholly owned venture of Lucent Technologies, makes visualization software for discovering patterns and relationships in huge databases. Over 300,000 data items appear as visual elements of different shapes, sizes, and colors, making obscure discoveries obvious.
This helps businesses save time and money by allowing people to quickly analyze large amounts of data and to use the resulting information to make better, more informed business decisions.

This demonstration uses the Visual Insights Developer's Toolkit, delivered as embeddable ActiveX components, that allows companies to easily incorporate powerful visual interfaces into their products. In this demonstration, the Visual Insights ActiveX components were embedded into a Microsoft Internet Explorer application.

Parallel Query Processing

Inter-query parallelism allows multiple queries to execute concurrently, resulting in a significant increase in throughput. DB2 UDB EEE incorporates advanced concurrency management mechanisms and workload balancing techniques to allow for extremely high throughput by utilizing all of the available resources.

Intra-query parallelism refers to the processing of parts of a single query at the same time using either intra-partition parallelism or inter-partition parallelism or both, resulting in dramatic query response time improvements. Intra-partition parallelism is the ability to take an individual task — such as a scan, sort, merge, or join — and generate multiple identical tasks operating on different pieces of data.
This speeds up a task within a particular partition. Inter-partition parallelism is breaking up a query into multiple parts across multiple partitions of a partitioned database. The query is performed in parallel across the partitions. DB2 UDB EEE generates a parallel execution strategy for all SQL statements using a cost-based relational database optimizer.
The parallel execution strategies in DB2 UDB EEE are executed asynchronously. Coordinator involvement is restricted to initializing slave tasks and to collecting final result sets.

Removing the Roadblocks to Affordable, High-Performance Clusters

The concept of clustering has been around for at least a decade. Large corporations have relied on clustering of big, proprietary computer systems to deliver the enterprise-class scalability, reliability, availability, and manageability they need for mission-critical applications, particularly in data warehousing and OLTP (online transaction processing). However, mainframe and RISC-based clustering solutions are based on proprietary, single-vendor architectures that are very expensive.

With the increased dependence on computer systems today, companies of all sizes are demanding affordable, high-performance clustering solutions. The use of data warehousing alone is expected to drive millions of databases into the multi-terabyte range. The companies participating in this demo and promoting VI Architecture are delivering the benefits of clustering to a wider audience than ever before, at lower price-points, and higher performance than ever before. They're doing it by building clusters around the common building blocks of Intel Architecture servers.

VI Architecture blends the economies of scale inherent in large-scale manufacturing with the computing power of high performance processors, offering IT managers attractive price/performance in a multi-system solution. VI Architecture makes it possible to scale a cluster by simply adding more nodes. This scalability enables a system to easily grow as business demands rise. Traditionally, if a company with a 4-way Windows NT Server node needed additional computing performance, the only way to achieve that higher level of performance

was to replace the server node with another, more powerful single node. VI Architecture's scalability protects a company's investment in their hardware and software infrastructure, because existing equipment and applications aren't forced into obsolescence.

VI Architecture also enhances the portability of applications by isolating the application from the details of hardware and operating system platforms. Rather than worrying about the idiosyncrasies of various operating systems and hardware driver interfaces, the developer can program to the VI Architecture interface, which is consistent across platforms. This not only increases the portability of the application but also simplifies the overall development process.

VI Architecture will enable a new class of scalable cluster products offering high performance, low total cost of ownership, and broad applicability. From a hardware product vendor's perspective, VI Architecture represents a framework for designing and building low-latency, high-reliability clusters of systems. From a software product vendor's perspective, the VI Architecture interface makes it easy to design applications capable of running on a variety of scalable, cost-efficient hardware platforms.

This demonstration proves that VI Architecture is a viable, formidable improvement to traditional clustering models. Companies of all sizes can now embrace affordable, flexible, commercially available clustering solutions that will give their businesses a new competitive edge.

CLARiiON Corporation
Coslin Drive
Southboro, MA 01772
USA
www.clariion.com

Dell Computer Corporation
One Dell Way
Round Rock, TX 78682
USA
www.dell.com

GigaNet
Concord Office Center
2352 Main Street
Suite 301
Concord, MA 01742
USA
www.giganet.com

IBM Corporation
One New Orchard Road
Armonk, NY 10504
USA
www.ibm.com

Intel Corporation
2200 Mission College Blvd.
PO Box 58119
Santa Clara, CA 95052-8119
USA
www.intel.com

qLogic Corporation
3545 Harbor Blvd.
Costa Mesa, CA 92626
USA
www.qlc.com

Visual Insights
263 Shuman Blvd.
PO Box 3050
Naperville, IL 60566-7050
USA
www.visualinsights.com

Printed in U.S.A.
09/98/12K

Intel Literature Center:
1-800-548-4725

Order No. 283683-001